

Hitting a Path: a Generalization of Weighted Connectivity via Game Theory

Dávid Szeszlér

Published online: 29 November 2018

Abstract Applying game-theoretical tools for measuring the reliability of a network has become very common. The basic idea is very natural: analyzing an appropriately defined attacker-defender game might give rise to a relevant security metric. In this paper we consider a very natural set of games: the Defender chooses a path P between two given nodes and the Attacker chooses a network element a (that is, an edge or a node). In all cases, the Attacker has to pay a given cost of attack $c(a)$; if, however, a is on P then he also gains a given profit of $d(a)$. We determine the value of various versions of this game and show that the thus arising reliability metrics provide a generalization of weighted connectivity of graphs. We also prove that the values of the games and optimum mixed strategies for both players can be computed in strongly polynomial time.

Keywords Network Reliability · Game Theory · Connectivity

Mathematics Subject Classification (2000) MSC 90B25 · MSC 91A80 · MSC 05C40

The research reported in this paper was supported by the Higher Education Excellence Program of the Ministry of Human Capacities in the frame of Artificial Intelligence research area of Budapest University of Technology and Economics (BME FIKP-MI/SC). Research supported by grants No. OTKA 108947 and 124171 of the Hungarian Scientific Research Fund.

Dávid Szeszlér
Department of Computer Science and Information Theory
Budapest University of Technology and Economics
Magyar Tudósok krt. 2., Budapest, 1117, Hungary
Tel.: +361-463-3162
Fax: +361-463-3157
E-mail: szeszler@cs.bme.hu

1 Introduction

Measuring the reliability of a network is one of the rich and complex areas of combinatorial optimization. Since the precise meaning of reliability highly depends on the application, a large variety of different reliability metrics have been proposed in the literature. Applying game-theoretical tools for measuring security has become very common. The basic idea is very natural: define a game between two virtual players, the Attacker and the Defender, such that the rules of the game capture the circumstances under which reliability is to be measured. Then analyzing the game might give rise to an appropriate security metric: the better the Attacker can do in the game, the lower the level of security is.

There is an abundance of recent books and papers on game-theoretical tools for measuring and increasing security. Since all aspects of security are obviously of utmost importance nowadays and game theory as a tool to address related problems presents itself very naturally, the literature on this topic is extremely diverse. Much of the arsenal of game theory has been employed on various applications which very often have little in common besides somehow being related to security. Interested readers are referred to the following books and surveys: [1], [9], [11], [12], [19].

In this paper, however, only the theory of two-player, zero-sum games, the simplest and probably most widely known subfield of game theory will be relied on to analyze a very naturally arising family of games and thus give rise to new graph reliability metrics that will turn out to be generalizations of weighted connectivity (meant in various ways).

The structure of this paper is as follows. We mention a few related results on network reliability games to motivate the topic of this paper below. In Section 2 we briefly summarize the necessary background in game theory. In Section 3 we define different versions of the *st*-Path Game and present the new contributions of the paper. Section 4 concludes the paper.

We follow the notation and terminology of [16]. In particular, for a function $f : E \rightarrow \mathbb{R}$ and a subset $U \subseteq E$, $f(U)$ denotes $\sum_{e \in U} f(e)$.

1.1 Network Reliability Games

As mentioned above, analyzing appropriately defined attacker-defender games is a natural approach for measuring the reliability of a network. Many such games known from the literature fall under the following framework.

Assume that an input graph G is given. G can either be directed or undirected depending on the application. Besides that, two weight functions are also part of the input: for each edge $e \in E(G)$ the “damage” caused by the loss of e (or in other words, the “importance” of e) is denoted by $d(e)$; furthermore, the cost of attacking an edge e is denoted by $c(e)$. Then a two-player, zero-sum game is defined on G between two virtual players, the Attacker and the Defender as follows. The Attacker chooses (or “attacks and destroys”) an edge

e of G . Simultaneously (or simply without knowing the Attacker's choice) the Defender chooses a subset of the edges $Z \subseteq E(G)$ that is thought of as some kind of "communication infrastructure" and the precise requirements on which vary in each application. Regardless of the Defender's choice, the Attacker has to pay the cost of attack $c(e)$ to the Defender. There is no further payoff if $e \notin Z$. If, on the other hand, $e \in Z$ then the Defender pays the Attacker the damage value $d(e)$.

Since these games are two-player, zero-sum games by definition, they have a unique Nash-equilibrium payoff V (which will simply be referred to as the *game value* in this paper) by Neumann's classic Minimax Theorem (see Section 2). Since V is the highest expected gain the Attacker can guarantee himself by an appropriately chosen mixed strategy, it makes sense to say that $\frac{1}{V}$ is a valid reliability metric in the sense corresponding to the specifics the game.

We remark that it might seem unrealistic in the above described framework that the Defender should receive the cost of attack $c(e)$ from the Attacker (as the Defender is indifferent to the costs and efforts associated with an attack, she is only affected by the damage caused). In other words, it would be more natural to assume that the above given payoffs only describe the Attacker's gain while the Defender's loss depends exclusively on e , Z and the damage value $d(e)$ (and is thus always bigger by $c(e)$ than the Attacker's gain). This would also imply that the game is not zero-sum any more. However, it is easily shown that the thus-obtained non-zero-sum game is essentially equivalent to the zero-sum game described above. This equivalency is due to the fact that the sum of the payoffs only depends on the choice of the Attacker and it more precisely means that Nash-equilibria of the two versions of the game are identical and the Attacker's Nash-equilibrium payoff is unique in the non-zero-sum version of the game and it is equal to the (unique) Nash-equilibrium payoff corresponding to the zero-sum version. (An analogous statement would not be true for the Defender.) The proof of this equivalency is a simple exercise (see [10, Lemma 1] for a proof). We will disregard this point in the remainder of the paper and focus on the zero-sum game versions described above.

The Spanning Tree Game Perhaps the most natural of the games falling under the above framework is the following. A connected, undirected graph G , a damage function $d : E(G) \rightarrow \mathbb{R}^+$ and a cost function $c : E(G) \rightarrow \mathbb{R}$ are given. The Attacker chooses an edge e of G and the Defender chooses a spanning tree T of G . Then the payoff from the Defender to the Attacker is $d(e) - c(e)$ if e is in T and $-c(e)$ otherwise. This game was first considered in the $d(e) \equiv 1$ and $c(e) \equiv 0$ case in [6] and then in the more general $d(e) \equiv 1$ (and $c(e)$ is arbitrary) case in [7]. For the general case, where $d(e) \geq 0$ and $c(e)$ are both arbitrary, the following was proved in [17]. (See Section 2 for a precise definition of the notion of game value.)

Theorem 1 *The game value of the Spanning Tree Game is*

$$\max_{\emptyset \neq U \subseteq E(G)} \frac{\text{comp}(G - U) - 1 - q(U)}{p(U)},$$

where $p(e) = \frac{1}{d(e)}$ and $q(e) = \frac{c(e)}{d(e)}$ for all $e \in E(G)$ and $\text{comp}(G - U)$ is the number of components of the graph obtained from G by deleting U . Furthermore, there exists a strongly polynomial algorithm to compute the game value of the Spanning Tree Game and an optimum mixed strategy for both players.

We remark that the above formula (without a corresponding strongly polynomial algorithm) was shown previously in the special case of $d(e) \equiv 1$ in [7]. Furthermore, in the special case of $c(e) \equiv 0$ the result of [16, Corollary 51.8a] is essentially equivalent to the above theorem in a non-game-theoretical setting. The running time of the algorithm given in [17] was later substantially improved to $O(|V|^4|E|^{\frac{1}{2}})$ in [2].

It also follows from the above theorem that the spanning tree game is capable of capturing a known graph reliability metric. The *strength* of a connected graph G was defined by Gusfield [8]. The idea is quite natural: if we remove a subset $U \subseteq E(G)$ of the edges then the efficiency of this “attack” against G can be measured by the ratio of the number of new components created and $|U|$ (that is, the “effort” required for the attack). Then it makes sense to define the reciprocal of the maximum efficiency of an attack to be a security metric: $\sigma(G) = \min \left\{ \frac{|U|}{\text{comp}(G-U)-1} : U \subseteq E(G), \text{comp}(G-U) > 1 \right\}$, where $\text{comp}(G-U)$ is the number of components of the graph obtained from G by deleting U . This notion was extended to a weighted version and its computability in strongly polynomial time was shown by Cunningham in [5]:

Definition 1 Assume that a connected graph G is given with a positive weight function $p : E(G) \rightarrow \mathbb{R}^+$ on its edges. Then

$$\sigma_p(G) = \min \left\{ \frac{p(U)}{\text{comp}(G-U)-1} : U \subseteq E(G), \text{comp}(G-U) > 1 \right\}$$

is called the *strength* of G with respect to p .

Corollary 1 ([17]) *The game value of the Spanning Tree Game is $\frac{1}{\sigma_p(G)}$ if $p(e) = \frac{1}{d(e)}$ and $c(e) = 0$ is assumed for all $e \in E(G)$.*

It is also worth mentioning that Theorem 1 was proved in [17] in a much more general, matroidal setting which gives rise to a number of natural extensions of the Spanning Tree Game and readily provides the corresponding modifications of the notion of graph strength. Interested readers are referred to [17] for the details.

The Arborescence Game A naturally arising, directed version of the Spanning Tree Game was defined in [17]. Call a subset of the nodes $R \subseteq V(G)$ of a digraph G a *source set* if every node of G is reachable from a node in R via a directed path. A vertex $r \in V(G)$ is a *source node* if $\{r\}$ is a single-element source set. Recall that an *arborescence* of a digraph G is a subset A of the arcs that is a spanning tree of the underlying undirected graph such that the digraph $(V(G), A)$ has a source node. (It is well-known and elementary that

the existence of an arborescence is equivalent to the existence of a source node.) Then the input of the Arborescence Game is a directed graph G that has a source node, a damage function $d : E(G) \rightarrow \mathbb{R}^+$ and a cost function $c : E(G) \rightarrow \mathbb{R}$. Analogously to the Spanning Tree Game, the Attacker chooses an arc e of G , the Defender chooses an arborescence A of G and the payoff from the Defender to the Attacker is $d(e) - c(e)$ if e is in A and $-c(e)$ otherwise.

Theorem 2 ([17]) *The game value of the Arborescence Game is*

$$\max_{\emptyset \neq U \subseteq E(G)} \frac{\text{source}(G - U) - 1 - q(U)}{p(U)},$$

where $p(e) = \frac{1}{d(e)}$ and $q(e) = \frac{c(e)}{d(e)}$ for all $e \in E(G)$ and $\text{source}(G - U)$ is the minimum cardinality of a source set in the digraph obtained from G by deleting U .

The existence of a strongly polynomial algorithm to compute the game value of the Arborescence Game and an optimum mixed strategy for both players was proved in [2].

Analogously to the undirected case, a directed version of the notion of graph strength presents itself very naturally and then the above theorem shows its connection to the Arborescence Game. The following notion was defined and its computability in strongly polynomial time was proved in [17]. The running time was then again improved in [2].

Definition 2 Assume that a directed graph G is given that has a source node; assume further that a positive weight function $p : E(G) \rightarrow \mathbb{R}^+$ is given. Then

$$\vec{\sigma}_p(G) = \min \left\{ \frac{p(U)}{\text{source}(G - U) - 1} : U \subseteq S, \text{source}(G - U) > 1 \right\}$$

is the *directed strength* of G with respect to p .

Corollary 2 ([17]) *The game value of the Arborescence Game is $\frac{1}{\vec{\sigma}_p(G)}$ if $p(e) = \frac{1}{d(e)}$ and $c(e) = 0$ is assumed for all $e \in E(G)$.*

For a survey of some more network reliability games of similar nature interested readers are referred to [18].

2 Preliminaries on Game Theory

A (*finite*) *two-player, zero-sum game* is given by a matrix M called the *payoff matrix*. Columns of M correspond to one of the players and rows of M to the other, so for the sake of simplicity one can refer to the two players as *Column Player* and *Row Player*. Columns and rows of M are called the *pure strategies* of the respective players. The matrix M defines the game in the following sense: both players choose one of their pure strategies (simultaneously, without

knowing each other's choices) and then the corresponding entry of M (that is, the one in the intersection of the chosen row and column) is paid by the Row Player to the Column Player. (Obviously, a negative payment means that in reality it is the Column Player who pays the absolute value of the amount to the Row Player.)

A *mixed strategy* of a player is a probability distribution on their pure strategies. If M is a $k \times n$ matrix then it is natural to store the Column Player's and the Row Player's mixed strategies as n -dimensional column vectors and k -dimensional row vectors, respectively. If we fix a pair of mixed strategies $\mathbf{x} \in \mathbb{R}^n$, $\mathbf{y} \in \mathbb{R}^k$ then the Column Player's expected gain (or, equivalently, the Row Player's expected loss) is obviously $\mathbf{y}M\mathbf{x}$. It is sensible for the Column Player to choose a mixed strategy \mathbf{x} that maximizes his worst case expected gain, therefore he is interested in finding an \mathbf{x} that maximizes the minimum value of $\mathbf{y}M\mathbf{x}$ over all possible mixed strategies \mathbf{y} of the Row Player; in other words, his job is $\max_{\mathbf{x}} \left\{ \min_{\mathbf{y}} \{ \mathbf{y}M\mathbf{x} \} \right\}$. Analogously, the Row Player's task is $\min_{\mathbf{y}} \left\{ \max_{\mathbf{x}} \{ \mathbf{y}M\mathbf{x} \} \right\}$; that is, she wants to minimize her worst case expected loss. Neumann's classic Minimax Theorem [14] states that these two values are equal for every payoff matrix M : $\max_{\mathbf{x}} \left\{ \min_{\mathbf{y}} \{ \mathbf{y}M\mathbf{x} \} \right\} = \min_{\mathbf{y}} \left\{ \max_{\mathbf{x}} \{ \mathbf{y}M\mathbf{x} \} \right\}$. This common value is called the *game value* corresponding to M . Since a pair of mixed strategies (\mathbf{x}, \mathbf{y}) that attain the corresponding optima is equivalent to the (more general) notion of a Nash-equilibrium in the special case of two-player, zero-sum games, the game value is also referred to as a (Nash-)equilibrium payoff in the literature (which is known to be unique in this special case). However, in this paper we will keep calling it the game value.

It is useful to mention that the description of the tasks of the two players can be simplified by observing that it is sufficient for a mixed strategy to "guard against" all pure strategies of the other player, that will imply that it also guards against all mixed strategies. For example, if every entry of the column vector $M\mathbf{x}$ is at least μ for a mixed strategy \mathbf{x} , that translates to saying that no matter which pure strategy the Row Player picks, the Column Player's expected gain is at least μ . However, this also implies $\mathbf{y}M\mathbf{x} \geq \mu$ for every mixed strategy \mathbf{y} (since $\mathbf{y}M\mathbf{x}$ is a convex combination of the entries of $M\mathbf{x}$). Hence the Column Player's task can also be described as maximizing the minimum entry of $M\mathbf{x}$ over all mixed strategies \mathbf{x} (and the Row Player's case is analogous).

The above also implies (as it is shown in many textbooks, see e.g. [13]) that two-player, zero-sum games are easy to handle algorithmically via linear programming: optimum mixed strategies for the game given by M can be found efficiently by solving the following linear program and its dual:

$$\max\{\mu : M\mathbf{x} \geq \mu \cdot \mathbf{1}, \mathbf{1} \cdot \mathbf{x} = 1, \mathbf{x} \geq \mathbf{0}\} \quad (1)$$

(where $\mathbf{1}$ denotes the all-1 vector). However, since the size of the payoff matrix M will be exponential in the size of the input graph G in all applications mentioned in this paper, this approach will not be viable.

3 The st -Path Game and its Variations

Motivated by the examples mentioned in the Introduction, the following definition might feel natural.

Definition 3 Assume that a directed graph G , two nodes $s, t \in V(G)$, a damage function $d : E(G) \rightarrow \mathbb{R}^+$ and a cost function $c : E(G) \rightarrow \mathbb{R}$ are given. Then the *Directed st -Path Game* is defined as follows: the Attacker chooses an edge e of G , the Defender chooses a directed path P from s to t (which is assumed to exist) and then the payoff from the Defender to the Attacker is $d(e) - c(e)$ if e is on P and $-c(e)$ otherwise. The *Undirected st -Path Game* is defined analogously with the single difference being that there G is undirected and the Defender chooses an undirected path between s and t .

Obviously, the above payoffs correspond to the framework described in the Introduction: the cost of attack $c(e)$ must be paid by the Attacker in all cases, but he receives the damage value $d(e)$ if he succeeds in hitting the st -path chosen by the Defender.

The Undirected st -Path Game was considered and solved in the $d(e) \equiv 1$ (and $c(e)$ is arbitrary) case in [3] and in the $c(e) \equiv 0$ (and $d(e)$ is arbitrary) case is [20]. (In [3], a generalization of the $d(e) \equiv 1$ case of the game was also solved: there the Attacker can target a subset of the edges of a given size and the Defender can choose two paths between two source-destination pairs.)

Before we claim the following theorem, we need to clarify some terminology: if G is a directed graph and $\emptyset \neq X \neq V(G)$ is a node set then the set of all edges leaving X for $V(G) - X$ is called a *cut*. If G is undirected then the set of all edges between X and $V(G) - X$ is called a *cut*. If $s \in X$ and $t \in V(G) - X$ for some nodes s and t then the cut is also referred to as an *st -cut*.

Theorem 3 *The game value of the Directed st -Path Game is*

$$\max \left\{ \left\{ \frac{1 - q(C)}{p(C)} : C \text{ is an } st\text{-cut} \right\} \cup \left\{ -c(e) : e \in E(G) \right\} \right\},$$

where $p(e) = \frac{1}{d(e)}$ and $q(e) = \frac{c(e)}{d(e)}$ for all $e \in E(G)$. Furthermore, the game value and an optimum mixed strategy for both players can be computed in strongly polynomial time.

Proof Let the value of the above maximum be μ . Assume first that $\mu = -c(e)$ for some $e \in E(G)$. Then if the Attacker targets e with a probability of 1, his total expected gain is obviously at least μ . Now assume that $\mu = \frac{1 - q(C)}{p(C)}$ for an st -cut C and let the Attacker use the following mixed strategy: assign a

probability of $\frac{p(e)}{p(C)}$ to every edge of C and 0 to the rest of the edges. Consider an arbitrary directed path P from s to t and fix an edge $e \in C$. Then e contributes to the Attacker's expected gain by $\frac{p(e)}{p(C)}(d(e) - c(e)) = \frac{1-q(e)}{p(C)}$ or $\frac{p(e)}{p(C)}(-c(e)) = \frac{-q(e)}{p(C)}$ depending on whether e is on P or not, respectively. Obviously, the contribution of edges $e \notin C$ is 0. Since the Attacker's expected gain is the total of the above contributions across all edges, this value is $\frac{T-q(C)}{p(C)}$, where $T = |C \cap E(P)|$. Observing that $T \geq 1$ follows from the fact that C is a cut, we get that the Attacker's total expected gain is at least $\frac{1-q(C)}{p(C)} = \mu$. Since in all cases the Attacker has a mixed strategy that guarantees him an expected gain of at least μ , the game value is also at least μ .

For every edge e of G let the capacity of e be $g(e) = \mu \cdot p(e) + q(e)$. Then, by the definition of μ , $g(e) \geq 0$ holds for every edge e and the total capacity of every st -cut in G is at least 1. (Indeed, the capacity of the cut C with respect to g is $g(C) = \mu \cdot p(C) + q(C)$. Therefore $g(C) \geq 1$ follows from $\mu \geq \frac{1-q(C)}{p(C)}$.) Therefore there exists a flow f from s to t of overall value 1 by the Ford-Fulkerson theorem.

It is well-known (see [16, Section 10.3]) that f is a non-negative linear combination of characteristic vectors of directed paths from s to t and directed cycles. Disregarding the directed cycles of such a decomposition we get that there exists a set of directed paths P_1, P_2, \dots, P_t in G from s to t and corresponding positive coefficients $\alpha_1, \alpha_2, \dots, \alpha_t$ such that $\sum_{i=1}^t \alpha_i = 1$ and

$$\sum_{i:e \in E(P_i)} \alpha_i \leq \mu p(e) + q(e) \quad (2)$$

holds for each edge e .

Now assume that the Defender uses the following mixed strategy: for every $1 \leq i \leq t$ she assigns the probability α_i to P_i and 0 to the rest of the st -paths. Then if the Attacker targets an edge e then her expected loss is

$$\begin{aligned} \sum_{i:e \in E(P_i)} \alpha_i (d(e) - c(e)) - \sum_{i:e \notin E(P_i)} \alpha_i c(e) &= \\ (d(e) - c(e)) \cdot \sum_{i:e \in E(P_i)} \alpha_i - c(e) \cdot \sum_{i:e \notin E(P_i)} \alpha_i &= \\ d(e) \cdot \sum_{i:e \in E(P_i)} \alpha_i - c(e) \leq d(e) (\mu p(e) + q(e)) - c(e) &= \mu \end{aligned}$$

by (2). Therefore this mixed strategy guarantees the Defender an expected loss of at most μ . Hence the game value is also at most μ , which implies by the above that it is exactly μ as claimed.

It also follows from the above that the game value is the minimum value of ν such that for the capacity function $g(e) = \nu \cdot p(e) + q(e)$ there exists a flow of overall value 1 and $g(e) \geq 0$ holds for every edge e . Indeed, $g(e) \geq 0$ implies $\nu \geq -c(e)$ and the existence of a flow of value 1 implies $\nu \cdot p(C) + q(C) \geq 1$ and

hence $\nu \geq \frac{1-q(C)}{p(C)}$ for every st -cut C . Therefore $\nu \geq \mu$ (where μ still denotes the maximum in the statement of the theorem). Furthermore, for the minimum of such ν 's obviously either $g(e) = 0$ must hold for some arc e or the maximum flow value must be exactly 1. In the first case $\nu = -c(e)$ for some arc e while in the second case there exists an st -cut C such that $\nu \cdot p(C) + q(C) = 1$ and hence $\nu = \frac{1-q(C)}{p(C)}$. In both cases we get $\nu = \mu$ as claimed.

Determining the minimum value of such a ν is obviously possible either by binary search or by linear programming, but these approaches would lead to weakly polynomial algorithms. On the other hand, this parametric flow problem was proved to be solvable in strongly polynomial time by parametric search in [4] or even more efficiently by $|E(G)|$ maximum flow computations in [15]. The latter algorithm is briefly described after the proof below. Once μ (that is, the minimum value of such a ν) is known, decomposing a flow f of overall value 1 with respect to the capacity function $\mu \cdot p + q$ into a convex linear combination of directed st -paths and directed cycles (which is obviously possible in strongly polynomial time) yields an optimum mixed strategy for the Defender according to the above. Furthermore, if $\mu = -c(e)$ for an arc e then targeting e with a probability of 1 is an optimum mixed strategy for the Attacker as shown above. If not then it also follows from the above that f is a maximum flow and hence one can easily compute an st -cut C of capacity 1 from f . C then yields an optimum mixed strategy for the Attacker as described in the first paragraph of the proof. \square

For the sake of completeness we briefly describe the strongly polynomial algorithm given in [15] for determining μ . As shown in the proof above, μ is the minimum value of ν such that $g(e) \geq 0$ holds for every edge e and there exists a flow of overall value 1 for the capacity function $g(e) = \nu \cdot p(e) + q(e)$. The algorithm keeps generating a sequence of values $\nu_0 < \nu_1 < \dots < \nu_t = \mu$ and corresponding flows f_0, f_1, \dots, f_t and cuts C_0, C_1, \dots, C_t such that $g_i(e) \geq 0$ always holds for the capacity function $g_i(e) = \nu_i \cdot p(e) + q(e)$ and f_i is a maximum flow and C_i is a minimum cut with respect to g_i . To initialize the algorithm, let $\nu_0 = \max\{-c(e) : e \in E(G)\}$; this choice already ensures $g_0(e) \geq 0$ for every edge e and $\nu_0 \leq \mu$ by the above theorem. Whenever, during the procedure, the value of ν_i is obtained for an $i \geq 0$, the algorithm computes a maximum flow f_i and a corresponding minimum cut C_i with respect to g_i . If the overall value of f_i is 1 then the algorithm terminates and outputs $\mu = \nu_i$. If not then it sets $\nu_{i+1} = \frac{1-q(C_i)}{p(C_i)}$ and continues the procedure (with ν_{i+1} instead of ν_i). Since $\nu_{i+1} \cdot p(C_i) + q(C_i) = 1$, the capacity of C_i with respect to g_{i+1} is 1 and hence the maximum flow with respect to g_{i+1} is at most 1. If, on the other hand, it is exactly 1, then $g_{i+1}(C) \geq 1$ for every cut C and $g_{i+1}(C_{i+1}) = 1$. In other words: $\nu_{i+1} \geq \frac{1-q(C)}{p(C)}$ for every cut C and $\nu_{i+1} = \frac{1-q(C_{i+1})}{p(C_{i+1})}$. This implies $\mu = \nu_{i+1}$ and hence the algorithm works correctly. It is also finite since the number of cuts is finite. The result proved in [15] implies that it terminates in at most $|E(G)|$ iterations and hence it yields a strongly polynomial algorithm.

We remark that the appearance of $\max\{-c(e) : e \in E(G)\}$ in the above theorem might look counterintuitive. However, its role can be supported by observing that it compensates for the non-monotonicity of the function $\frac{1-q(Z)}{p(Z)}$ in the following sense: denote the maximum of $\frac{1-q(Z)}{p(Z)}$ across all subsets $Z \subseteq E(G)$ that contain an st -cut as a subset by μ' . Then μ' can be strictly bigger than if this maximum were only taken across st -cuts. On the other hand, it is easy to show that μ' is less than or equal to (and can be strictly less than) the maximum in Theorem 3. (Indeed, since the argument of the first paragraph of the proof of Theorem 3 can be applied on edge sets containing an st -cut, it follows that the Attacker can guarantee himself an expected gain of μ' which implies $\mu' \leq \mu$.)

An analogous theorem holds for the Undirected st -Path Game:

Theorem 4 *The game value of the Undirected st -Path Game is*

$$\max \left\{ \left\{ \frac{1-q(C)}{p(C)} : C \text{ is an } st\text{-cut} \right\} \cup \left\{ -c(e) : e \in E(G) \right\} \right\},$$

where $p(e) = \frac{1}{d(e)}$ and $q(e) = \frac{c(e)}{d(e)}$ for all $e \in E(G)$. Furthermore, the game value and an optimum mixed strategy for both players can be computed in strongly polynomial time.

Proof Let the value of the above maximum be μ . Replace every edge $e = \{u, v\}$ of G by the directed arcs $e' = \vec{uv}$ and $e'' = \vec{vu}$ and denote the obtained digraph by D . Let $d(e') = d(e'') = d(e)$ and $c(e') = c(e'') = c(e)$ for every edge. Applying Theorem 3 and noting that the total capacity of each st -cut of D is equal to the total capacity of the corresponding st -cut of G , we get that the game value of the Directed st -Path Game played on D is μ . Consequently, there exists a pair of mixed strategies $\mathbf{x}_D : E(D) \rightarrow [0, 1]$ and $\mathbf{y}_D : \mathcal{P}_D \rightarrow [0, 1]$ such that \mathbf{x}_D guarantees the Attacker an expected gain of μ and \mathbf{y}_D guarantees the Defender an expected loss of μ in the Directed st -Path Game played on D , where \mathcal{P}_D denotes the set of directed paths from s to t in D .

Now let $\mathbf{x}_G(e) = \mathbf{x}_D(e') + \mathbf{x}_D(e'')$ for every $e \in E(G)$. Since the payoff corresponding to a path P and an edge e in the game played on G is at least as big as the payoff corresponding to the directed version of P and either e' or e'' in the game played on D , it follows that \mathbf{x}_G guarantees the Attacker of the game played on G an expected gain of at least μ . Hence the game value is also at least μ .

We claim that it can be assumed without loss of generality that for every $e \in E(G)$ at most one of e' and e'' is contained in a directed path P for which $\mathbf{y}_D(P) > 0$. Assume to the contrary that $e' = \vec{uv}$ is contained in the path P_1 and $e'' = \vec{vu}$ is contained in the path P_2 and $0 < \mathbf{y}_D(P_1) \leq \mathbf{y}_D(P_2)$. Let $P_{1,2}$ be a directed path from s to t that is contained in the walk consisting of the first part of P_1 leading from s to u and the second part of P_2 leading from u to t . Define $P_{2,1}$ analogously. Now decrease $\mathbf{y}_D(P_1)$ and $\mathbf{y}_D(P_2)$ by $\mathbf{y}_D(P_1)$ and increase $\mathbf{y}_D(P_{1,2})$ and $\mathbf{y}_D(P_{2,1})$ by $\mathbf{y}_D(P_1)$. Since $\sum\{\mathbf{y}_D(P) : e \text{ is on } P\}$

was not increased for any arc e of D , the modified \mathbf{y}_D still guarantees an expected loss of at most μ to the Defender, however, the number of paths P with a positive $\mathbf{y}_D(P)$ containing e' or e'' got decreased. Therefore repeating this modification as many times as needed yields the claim.

Now for every st -path P of G let $\mathbf{y}_G(P) = \mathbf{y}_D(P')$, where P' denotes the directed version of P . Then the claim of the above paragraph implies that \mathbf{y}_G guarantees the Defender an expected loss of at most μ in the game played on G . Hence the game value is also at most μ .

Since \mathbf{x}_D and \mathbf{y}_D can be computed in strongly polynomial time by Theorem 3 and \mathbf{x}_G and \mathbf{y}_G can easily be computed from these according to the above, the proof is complete. \square

Let $p : E(G) \rightarrow \mathbb{R}^+$ be a positive valued weight function on the edges of G . Recall that $\lambda_p(s, t)$, the *weighted edge-connectivity between s and t with respect to p* is the minimum value of $p(C)$, where C is an st -cut (both in the directed and the undirected case). It follows directly from Theorems 3 and 4 that the (Directed and Undirected, respectively) st -Path Game is capable of capturing the notion of $\lambda_p(s, t)$ and its game value can be considered as a sensible generalization of (the reciprocal of) $\lambda_p(s, t)$. The following was also proved in [20] (for the undirected case).

Corollary 3 *If $c(e) = 0$ for all $e \in E(G)$ then the game value of (both the Directed and the Undirected) st -Path Game is $\frac{1}{\lambda_p(s, t)}$, where $p(e) = \frac{1}{d(e)}$ for every edge e .*

We remark that while the reciprocal of the game value of the st -Path Game generalizes the notion of weighted edge-connectivity between two given nodes, it can easily be modified so as to capture the notion of (general) weighted edge-connectivity $\lambda_p(G)$, that is, the minimum value of $p(C)$ taken across all cuts of G . Indeed, consider the following modification of the game: first the Attacker chooses two distinct nodes $s, t \in V(G)$ and declares them to the Defender. Then the Defender chooses a path P from s to t and (simultaneously) the Attacker chooses an edge e of G . Finally, the payoff from the Defender to the Attacker is the same: $d(e) - c(e)$ or $-c(e)$ depending on whether e is on P or not, respectively. It follows easily from Theorems 3 and 4 that the value of this game is

$$\mu = \max \left\{ \left\{ \frac{1 - q(C)}{p(C)} : C \text{ is a cut of } G \right\} \cup \left\{ -c(e) : e \in E(G) \right\} \right\},$$

where $p(e) = \frac{1}{d(e)}$ and $q(e) = \frac{c(e)}{d(e)}$ for all $e \in E(G)$ (both in the directed and the undirected case). To show this, the Attacker should first choose a cut C for which $\frac{1 - q(C)}{p(C)}$ is maximum and choose $s \in X$ and $t \notin X$ arbitrarily, where C consists of the edges leaving X . Since the value of the thus arising st -Path Game is μ according to Theorems 3 and 4, it follows that the Attacker can guarantee himself an expected gain of at least μ and the Defender can

guarantee herself an expected loss of at most μ and hence the value of this modified game is indeed μ . It also follows directly that if $c(e) = 0$ is assumed for all $e \in E(G)$ then the game value is $\frac{1}{\lambda_p(G)}$.

While all the above results offered a way of capturing and generalizing the notion of weighted edge-connectivity, it is also possible to modify the *st*-Path Game so as to achieve a similar connection with weighted node-connectivity:

Definition 4 Assume that a (directed or undirected) graph G and two nodes $s, t \in V(G)$ are given such that t is reachable from s in G via a (directed or undirected) path, but there is no direct edge from s to t . Furthermore, a damage function $d : V_0 \rightarrow \mathbb{R}^+$ and a cost function $c : V_0 \rightarrow \mathbb{R}$ are also given, where $V_0 = V(G) \setminus \{s, t\}$. Then the *Node st-Path Game* is defined as follows: the Attacker chooses a vertex $v \in V_0$, the Defender chooses a (directed or undirected) path P from s to t and then the payoff from the Defender to the Attacker is $d(v) - c(v)$ if v is on P and $-c(v)$ otherwise.

A set of nodes $Z \subseteq V_0$ is said to *cover all st-paths* if every path from s to t contains at least one node in Z .

Theorem 5 *The game value of the Node st-Path Game is*

$$\max \left\{ \left\{ \frac{1 - q(Z)}{p(Z)} : Z \subseteq V_0, Z \text{ covers all } st\text{-paths} \right\} \cup \left\{ -c(v) : v \in V_0 \right\} \right\},$$

where $p(v) = \frac{1}{d(v)}$ and $q(v) = \frac{c(v)}{d(v)}$ for all $v \in V_0$. Furthermore, the game value and an optimum mixed strategy for both players can be computed in strongly polynomial time.

Proof Assume first that G is directed. Let the value of the above maximum be μ and let $w \in V_0$ such that $c(w) = \min\{c(v) : v \in V_0\}$. Split every node $v \in V_0$ into two nodes: replace v by v_1 and v_2 and let $s_1 = s_2 = s$ and $t_1 = t_2 = t$ for the sake of consistency of the notation. Now add the new arc $e_v = \overrightarrow{v_1 v_2}$ for every $v \in V_0$ and replace every $e = \overrightarrow{u v} \in E(G)$ by $e' = \overrightarrow{u_2 v_1}$. Denote the obtained digraph by D and let $c(e_v) = c(v)$ and $d(e_v) = d(v)$ for every $v \in V_0$ and let $c(e') = c(w) + 1$ and $d(e') = 1$ for every $e \in E(G)$.

Consider an *st*-cut C of D . Using the observation that $\frac{q(X)}{p(X)} \geq \min\{c(v) : v \in X\}$ holds for every $\emptyset \neq X \subseteq V(G)$ one easily checks that $\frac{1 - q(C)}{p(C)} \leq -c(w)$ holds if $e' \in C$ for an edge $e \in E(G)$ (that is, if C contains an arc that corresponds to an original arc of G). This implies that the maximum in Theorem 3 computed for D is either $-c(w) = -c(e_w)$ or it is attained by an *st*-cut C of D that only contains arcs corresponding to nodes of G . Consider the latter case and let the corresponding node set of G be Z (that is, $C = \{e_v : v \in Z\}$). It is again easy to check that C is an *st*-cut of D if and only if Z covers all paths from s to t in G . All these together imply by Theorem 3 that the value of the Directed *st*-Path Game played on D is μ .

Consequently, there exists a pair of mixed strategies $\mathbf{x}_D : E(D) \rightarrow [0, 1]$ and $\mathbf{y}_D : \mathcal{P}_D \rightarrow [0, 1]$ such that \mathbf{x}_D guarantees the Attacker an expected gain

of μ and \mathbf{y}_D guarantees the Defender an expected loss of μ in the Directed st -Path Game played on D , where \mathcal{P}_D denotes the set of directed paths from s to t in D .

Now let $\mathbf{x}_G(v) = \mathbf{x}_D(e_v)$ for every $v \in V_0 \setminus \{w\}$ and let $\mathbf{x}_G(w) = \mathbf{x}_D(e_w) + \sum\{\mathbf{x}_D(e') : e \in E(G)\}$. Observe that for every $e \in E(G)$ the payoff corresponding to a path P and e' is always less than or equal to the payoff corresponding to P and e_w . (Indeed, depending on whether e' is on P or not, the payoff corresponding to P and e' is either $d(e') - c(e') = -c(w)$ or $-c(e') = -c(w) - 1$ and hence it is at most $-c(w)$. Similarly, the payoff corresponding to P and e_w is either $d(w) - c(w)$ or $-c(w)$, so it is at least $-c(w)$.) This implies that \mathbf{x}_G guarantees the Attacker an expected gain of at least μ in the Node st -Path Game played on G and thus the game value is at least μ .

On the other hand, let $\mathbf{y}_G(P) = \mathbf{y}_D(P')$ for every st -path of G , where P' is the st -path in D corresponding to P . It follows directly that \mathbf{y}_G guarantees the Defender an expected loss of μ in the Node st -Path Game played on G . Hence the game value is also at most μ and thus it is exactly μ .

Obviously, \mathbf{x}_G and \mathbf{y}_G can be computed in strongly polynomial time since the same holds for \mathbf{x}_D and \mathbf{y}_D by Theorem 3 and \mathbf{x}_G and \mathbf{y}_G can easily be computed from these according to the above.

Finally, assume that G is undirected. Replace every edge $e = \{u, v\}$ of G by the directed arcs \vec{uv} and \vec{vu} and denote the obtained digraph by G' . Since directed st -paths in G' directly correspond to undirected st -paths in G and $V(G') = V(G)$, it follows that the Node st -Path Games played on G and G' are identical (including that the payoffs are also equal in each case). Since the theorem was already proved to hold on G' , observing that node sets covering st -paths in G and G' are also identical yields the theorem on G too. \square

Obviously, the node splitting technique applied in the above proof could also be used to handle the version of the st -Path Game where both edges and nodes of a graph have given cost and damage values and the Attacker is allowed to choose either an edge or a node (different from s and t), however, the details of this are omitted here.

Recall that for any positive valued weight function $p : V(G) \rightarrow \mathbb{R}^+$, the *weighted node-connectivity between s and t with respect to p* is the minimum value $\kappa_p(s, t)$ of $p(Z)$, where Z ranges across all node sets covering all st -paths (both in the directed and the undirected case). The following is implied directly by Theorem 5:

Corollary 4 *If $c(v) = 0$ for all $v \in V_0$ then the game value of (both the Directed and the Undirected) Node st -Path Game is $\frac{1}{\kappa_p(s, t)}$, where $p(v) = \frac{1}{d(v)}$ for every $v \in V_0$.*

Analogously to what was remarked after Corollary 3, one can modify the rules of the Node st -Path Game (by letting the Attacker choose s and t first) so as to capture and generalize the notion of general weighted node-connectivity of a graph, we omit the details.

4 Conclusions

In this paper we defined and analyzed different versions of a two-player, zero-sum game played on a graph by two virtual players, the Defender and the Attacker. In all versions of the game the Attacker's goal was to hit a path chosen by the Defender between two given nodes. We determined the game value of all versions of the game and showed that this, as well as optimum mixed strategies for both players, can be computed in strongly polynomial time. This also implied that in the special case where the cost of attack for the Attacker is zero, the game value always coincides with the reciprocal of weighted connectivity understood in the sense corresponding to the specific version of the game. This observation led us to the conclusion that the games discussed in the paper are capable of capturing the notion of weighted connectivity of a graph and their game values can be considered as a sensible generalization of (the reciprocal of) weighted connectivity.

References

1. T. Alpcan and T. Başar, *Network security: A decision and game-theoretic approach*, Cambridge University Press (2010).
2. M. Baiou and F. Barahona, *Faster Algorithms for Security Games on Matroids*, *Algorithmica* (2018), online version.
3. G. Calinescu, S. Kapoor, M. Quinn and J. Shin, *Adversary Games in Secure/Reliable Network Routing*, *Proc. of Game Theory for Networks: Second International ICST Conference, GameNets 2011*, Shanghai, China, pp. 249-264, Springer Berlin Heidelberg (2012).
4. E. Cohen and N. Megiddo, *Algorithms and complexity analysis for some flow problems*, *Algorithmica*, vol. 11(3), pp. 320-340 (1994).
5. W. H. Cunningham, *Optimal attack and reinforcement of a network*, *Journal of the ACM (JACM)*, vol. 32(3), pp. 549-561 (1985).
6. A. Gueye, J. C. Walrand and V. Anantharam, *Design of network topology in an adversarial environment*, *Proc. of the 1st International Conference on Decision and Game Theory for Security, GameSec10*, Berlin, Germany, pp. 1-20 (2010).
7. A. Gueye, J. C. Walrand and V. Anantharam, *A network topology design game: How to choose communication links in an adversarial environment*, *Proc. of the 2nd International ICST Conference on Game Theory for Networks, GameNets*, vol. 11 (2011).
8. D. Gusfield, *Connectivity and edge-disjoint spanning trees*, *Information Processing Letters*, vol. 16.2, pp. 87-89 (1983).
9. Z. Han, *Game theory in wireless and communication networks: theory, models, and applications*, Cambridge University Press (2012).
10. A. Laszka and D. Szeszlér, *Hide and Seek in Digital Communication: The Steganography Game*, *Proc. 9th Hungarian-Japanese Symposium on Discrete Mathematics and Its Applications*, Fukuoka, Japan, pp. 126-136 (2015).
11. R. Machado and S. Tekinay, *A survey of game-theoretic approaches in wireless sensor networks*, *Computer Networks*, vol. 52(16), pp 3047-3061 (2008).
12. M. H. Manshaei, Q. Zhu, T. Alpcan, T. Bacsar and J. P. Hubaux, *Game theory meets network security and privacy*, *ACM Computing Surveys (CSUR)*, vol. 45.3 , Article No. 25 (2013).
13. J. Matoušek, B. Gärtner, *Understanding and Using Linear Programming*, 226 pages. Springer, Berlin, Heidelberg (2007).
14. J. v. Neumann, *Zur Theorie der Gesellschaftsspiele*, *Mathematische Annalen*, vol. 100(1), pp. 295-320 (1928).

15. T. Radzik, Parametric Flows, Weighted Means of Cuts and Fractional Combinatorial Optimization, in: Complexity in Numerical Optimization, ed. P. Pardalos, World Scientific Co., pp. 351-386 (1993).
16. A. Schrijver, Combinatorial Optimization: Polyhedra and Efficiency, Algorithms and Combinatorics vol. 24, 1879 pages. Springer, Berlin, Heidelberg (2003).
17. D. Szeszlér, Security games on matroids, Mathematical Programming vol. 161(1), pp. 347-364 (2017).
18. D. Szeszlér, Measuring Graph Robustness via Game Theory, Proc. 10th Japanese-Hungarian Symposium on Discrete Mathematics and Its Applications, Budapest, pp. 473-482 (2017).
19. M. Tambe, Security and game theory: algorithms, deployed systems, lessons learned. Cambridge University Press, New York (2012).
20. A. Washburn and K. Wood, Two-Person Zero-Sum Games for Network Interdiction, Operations Research vol. 43(2), pp. 243-251 (1995).